Information Architecture Targets Infrastructure Services

When people talk about client/server applications, you'll frequently hear much discussion about which client platform(s) to support, which DBMS (database management system) to use for the server, and issues like functionality vs. interoperability. If the conversation turns to the Internet/WorldWide Web, you're likely to hear a lot about Mosaic vs. Netscape, HTML II vs. HTML Plus, the White House home page, and other such topics. For information sharing, the conversation might address which desktop applications to use, how to educate people about available information, who has legitimate reason to access which information, etc.

Throughout these discussions, however, one question frequently overlooked is "What about the connection?" After all, even the best client/server application is only as good as the connection between the client and server, even the best Internet browser is no faster than the connection, and information sharing without the connection is no more effective than interoffice mail.

In recognition of this, one of the three main recommendations from Phase I of the Information Architecture (IA) project was that "The Laboratory Leadership Council endorse an institutional information infrastructure." When the LLC unanimously approved this recommendation in May 1994, the stage was set for the formation of the IA Infrastructure Team.

As the basic connecting link among the various components of the IA, the infrastructure's relationship to the data, desktop, applications, and data warehouse components can be modeled as shown in Figure 1.

For further discussion concerning roles of the Data, Desktop, Applications, and Data Warehouse IA Teams, see the November 1994 issue of BITS.

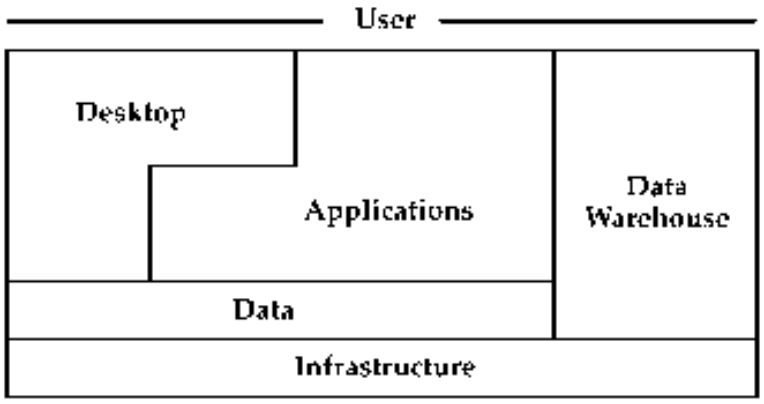

Figure 1. **Basic relationship of IA teams**

More than Just a Wire

As the Infrastructure Team began its work, it soon learned that "connection" means a good deal more than just "wire." As shown in the following table, the scope of the infrastructure also includes protocols and security, and it leads quickly into areas such as services and training.

As further explained below, many of the specific subject areas overlap with work already being done by others. The table shows that the subject areas reach into network management procedures and tools, which in turn extend the infrastructure all the way to the user's desktop. A revised model of the teams' relationships is shown in Figure 2.

General Areas	Specific Subjects
Foundations for Connectivity	LANLnet Topology Network Management Network Protocols
Basic Services	Newsgroups Automated Backups and Recovery E-mail Internet/WorldWide Web
Database Access	On-line Documents Distributed File Systems Naming Services
Security	Security Architecture Encryption, Authorization, Authentication Security Testing
Training	User Education Awareness Training

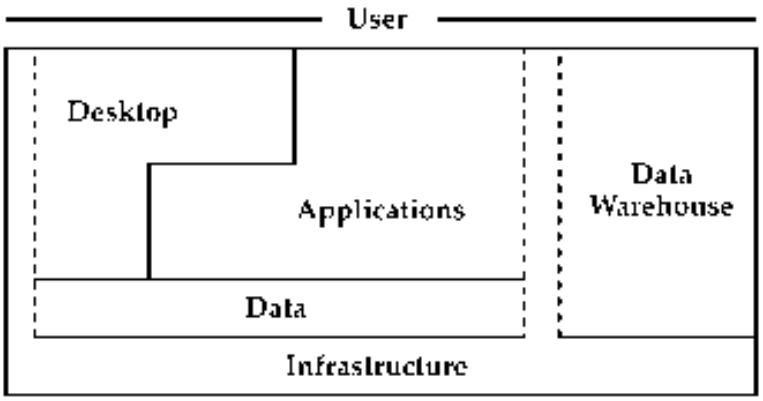

Figure 2. Revised Relationship of IA Teams

In this revised model, the dotted lines emphasize the overlap between the infrastructure and other teams—areas of shared interest where the various teams must actively collaborate. Indeed, collaboration is a common theme for the infrastructure team, both within and outside the IA project itself.

As Allen Mathews (DX-13), team leader for the IA Infrastructure Team, says, "Fortunately, many of the topics we are concerned with are already being addressed by Laboratory-sponsored projects," such as LANLnet, the CIC-2 Hypermedia Team, and the Computer Security Working Group.

"In these instances," Mathews continues, "our goal is to leverage off work already being done by others and help to make the result consistent with the rest of the IA."

Selected Subject Areas

To handle the breadth of the scope, the Infrastructure Team is chartering separate teams to address each of the specific subject areas. Focusing on just a few of these areas can demonstrate how the team coordinates its work with others and how the results will impact the Laboratory community.

Network Management: A primary objective of the team is to develop, in conjunction with Laboratory network managers, a consistent set of guidelines for network management. The fundamental purpose here is to iron out inconsistencies that become barriers to providing fundamental services, implementing distributed processing, and navigating through the various areas of the Laboratory's communications infrastructure, regardless of underlying protocols or network software (e.g., users shouldn't have to care whether a network is running Novell or AppleTalk).

Encryption, Authorization, & Authentication: Currently, too many applications are unable to move forward because of cumbersome security requirements. Private industry, faced with the pressure to bring commerce to the Internet, is rapidly advancing in its techniques for protecting bank accounts, credit card numbers, and similar sensitive information. By exploiting these techniques, the IA hopes to move security to the data level, open up applications to wider use, enable more effective sharing of information, and foster the trend toward a unified underlying infrastructure (as opposed to the current need to split the open/administrative and secure networks).

Automated Backups and Recovery: The information on your local hard disk is a Laboratory asset. Laboratory resources have been expended to create and maintain that information, and it is in everyone's interest that that information be protected from loss. One service that can not only achieve this but also ease your system administration work load is automated backups—just leave your machine on at night and the infrastructure will reach out and backup your disk to a central location. Then, if your machine fails (or you accidentally "rm-r" in the wrong directory) you can easily recover the lost information. The IA role in developing this capability is to develop the guidelines for when and how it can be effectively and responsibly implemented.

Internet/WorldWide Web: There is a revolution underway in our methods of accessing and using information. While the government might talk of the National Information Infrastructure (the superhighway we look forward to) and the United Nations might talk about the promise of its Open Systems Interconnection (the X. protocols), universities and private industry have continued to push ahead with the development of the Internet. A series of robust, well-tested protocols is already out there, and the server/browser technologies are rapidly evolving (if your browser is more than three months old, it's probably ready for an upgrade). Along with the availability of these tools comes a wide range of problems such as

IA Infrastructure Team Mission Statement

Mission: Lead the Laboratory in defining, developing, and promoting guidelines, standards, and procedures for the information infrastructure at Los Alamos.

Definition: The information infrastructure comprises the connection, equipment, software, services, security, protocols, management, and consistent support that provide and maintain the connections among Laboratory citizens' and external stakeholders' desktops, central computers, and other information resources, both local and remote.

how to maintain links to legacy methods (Gopher) while exploiting the potential of new tools (Mosaic, Netscape, etc.); how to present information in a radically new format (it doesn't look like a page, read like a page, or work like a page); how to find what you're looking for and make it easy for others to find what you've put out there; and what is responsible use.

Awareness Training: Many of the information resources that are already available are under used. This phenomenon is likely to get even more pronounced in the future as the rate of innovation continues to increase. The fundamental goal here is partially to let people know what is available but even more importantly to develop the flexible mechanisms that can adapt easily and gracefully to future needs.

The subject areas listed above are but a portion of what the IA Infrastructure Team has targeted, but they should illustrate the pervasive impact that infrastructure guidelines and standards will have on the Laboratory community.

Participation is Welcome

A cornerstone of the IA project as a whole is its emphasis on consensus building among as wide a sample of the Laboratory community as possible. Although the IA project has been empowered by the LLC to develop guidelines and standards, it is the active involvement of Laboratory citizens—scientific, engineering, and administrative alike—that ensures the deliverables will meet the real world needs of those citizens. Hence, the project welcomes participation in a number of ways:

• Anyone who wants to is welcome to join an IA team; in the history of the project, no one has ever been turned away.

• Groups are continually forming to address specific subject areas, and you are welcome to work on any subject areas that interest you. (See sidebar for details on how to find information.)

• If you don't have time to actively serve on a team, you can still follow the team's progress by subscribing to its e-mail distribution list.

• The project periodically holds forums to discuss its progress with anyone who chooses to attend; these forums are especially valuable to the project because they tend to reveal new perspectives.

• Before an idea becomes a guideline, the IA project issues requests for comment (RFC); anyone can respond to these requests, and all responses will be taken seriously.

• Before the idea, along with input from its RFCs, can evolve into a standard, it is issued as a guideline that Laboratory citizens are encouraged but not required to adhere to; this is a test period, and any input about the effects of the guideline, good or bad, is welcome.

As Mathews says, "We have set an ambitious schedule. Anyone at the Laboratory who is interested in helping us on these topics is welcome to join our team."

To join the IA Infrastructure Team, contact Allen Mathews (DX-13, 7-9055, arm@lanl.gov). For instructions on joining other IA teams, see the November 1994 issue of BITS or look beneath the IA on-line home page under "Information Architecture Teams" (see accompanying box for instructions on reaching the IA Gopher/Mosaic home page).

Tad Lane, tad@lanl.gov,
(505) 667-0886
Communications Arts and Services (CIC-1)

How to Track IA Progress

To follow the project's progress on the LANL Gopher/Mosaic server, look under "Computing at LANL/Information Architecture Project." The URL is

http://www.lanl.gov:8000/info-arch

The IA project currently maintains the following open e-mail distribution lists for its teams:

• Applications Team: ia-app@lanl.gov
• Data Team: ia-data@lanl.gov
• Data Warehouse Team: ia-dataw@lanl.gov
• Desktop Team: ia-desk@lanl.gov
• Infrastructure Team: ia-inf@lanl.gov

Additional lists form and dissolve as work progresses on specific areas. For an up-to-date listing, look beneath the IA home page under "Introduction to Information Architecture Project."

To follow or join newsgroup discussions of RFCs and proposed guidelines and standards, look on the LANL Gopher/Mosaic server under "Network News (USENET)/LANL Specific Newsgroups/info-arch." The URL is

gopher://infoserver.lanl.gov:4320/1nntp%20ls%20lanl.info-arch%20

Related non-IA e-mail distribution lists of interest include:

• Internet Information Systems User Group: iisug@lanl.gov
• WorldWide Web User Group: wwwug@lanl.gov
• LANLnet Distribution List: lanlnet@lanl.gov